

Leveraging Covariates in Regression Discontinuity Designs*

Matias D. Cattaneo[†] Filippo Palomba[‡]

January 16, 2025

Keywords

Causal inference, treatment effect estimation, regression discontinuity, covariate adjustment, head start.

Teaser

Proper use of covariates in Regression Discontinuity designs can enhance empirical scientific discoveries and evidence-based policy decisions.

Elevator Pitch


It is common practice to incorporate additional covariates in empirical economics. In the context of Regression Discontinuity (RD) designs, covariate adjustment plays multiple roles, making it essential to understand its impact on analysis and conclusions. Typically implemented via local least squares regressions, covariate adjustment can serve three main distinct purposes: (i) improving the efficiency of RD average causal effect estimators, (ii) learning about heterogeneous RD policy effects, and (iii) changing the RD parameter of interest. Unfortunately, “fixing” a “broken” RD design through covariate adjustment often requires implausibly strong assumptions about the econometric model.

*Invited contribution prepared for the *IZA World of Labor*.

[†]Department of Operations Research and Financial Engineering, Princeton University.

[‡]Department of Economics, Princeton University.

Graphical Abstract



Key Findings

- **Pros**

- + Precision of RD causal effect estimators can be improved by correctly using pre-intervention covariates.
- + When properly implemented, covariate adjustment can uncover interesting heterogeneous RD causal effects.
- + Various other RD policy effects of interest can be learned by correctly incorporating covariates in the analysis.

- **Cons**

- Understanding empirical results based on covariate adjustment RD estimation requires careful consideration.
- Covariate adjustment cannot restore the validity of an RD design without strong parametric assumptions.
- Empirical work leveraging covariates in RD designs is often undisciplined and ad-hoc, potentially leading to invalid empirical findings and policy prescriptions.

Authors' main message

When applied correctly, covariate adjustment in RD designs can significantly enhance empirical analysis and strengthen policy conclusions. However, the use of covariates across RD studies is often inconsistent and ad-hoc, undermining both the credibility and replicability of findings. Adopting best methodological practices for covariate adjustment in RD designs can improve efficiency and support rigorous heterogeneity analysis. Additionally, pre-intervention covariates can be leveraged to modify the RD parameter of interest, though this requires adopting additional, sometimes stringent, modeling assumptions.

Motivation

In program evaluation, a primary goal is to determine the causal effect of a policy on an outcome. Among the numerous experimental and non-experimental research designs available in the literature (see [1] and references therein), the RD design is widely used in empirical economics. A landmark example of its application is found in [30], which examines the causal impact of the U.S. federal anti-poverty program Head Start on child mortality. The study leveraged a discontinuity in program funding that occurred in 1965, when Head Start was first implemented. To ensure adequate representation of the nation's poorest communities in a nationwide grant competition for program funds, the federal government provided assistance to the 300 poorest U.S. counties to help them prepare and submit funding applications. This support resulted in a surge of participation and funding rates among these counties, creating a discontinuity at the 300th poorest county, and thus enabled the use of an RD design for policy evaluation. Using flexible parametric methods, [30] found that access to increased Head Start funding led to a significant reduction in child mortality rates between 1973 and 1983 among children aged five to nine, specifically for causes impacted by the program's health services. The reported decrease in mortality was substantial, from approximately 3.2 to 1.9 deaths per 100,000 children. More recently, [19] re-examined the Head Start data using modern RD methods (see [18] and references therein), and showed that the reduction in child mortality is robust to different assumptions and estimation strategies.

The Head Start application serves as a prototypical empirical example of the canonical (sharp) RD design, in which each unit is assigned a score, and treatment is given to all units whose score exceeds a fixed cutoff, while those below the cutoff receive no treatment. Specifically, in the Head Start context, the units are counties, the score is the poverty index, and the cutoff is the poverty level of the 300th poorest U.S. county (approximately equal to 59.2). The core identification strategy hinges on the sudden shift in the likelihood of receiving treatment at the cutoff, which allows for comparing units just above and below the threshold to estimate the causal effect of the treatment. In practice, local least squares linear regression employing only those units whose scores are close to the cutoff point determining policy assignment is often used to estimate the RD average causal effect at the cutoff (see [13] and [14] for comprehensive practical introduction).

The Head Start data employed by [30] also included pre-intervention covariates and other post-intervention outcomes. The authors, and later [19], utilized these variables for falsification of the RD design: it is common practice to verify that the RD causal effect is statistically indistinguishable from zero when the outcome variable is known to be unaffected by the intervention. Furthermore, pre-intervention covariates can be leveraged to improve the precision of estimating the average causal effect at the cutoff, as well as to explore treatment effect heterogeneity. In this article, we discuss the role of covariate adjustment in RD designs and demonstrate their application in the context of the Head Start study. Following [9], we leverage nine pre-intervention covariates: *population*, *population between age 14 and 17*, *population between age 5 and 34*, *population with age 25+*, *% attending school between age 14 and 17*, *% attending school between age 5 and 34*, *% of urban areas*, *% of urban areas*, and *% of Black population*. The eight covariates excluding *population* are only used for improving estimation efficiency, and thus we refer to them as “covariates for efficiency”. The pre-intervention covariate *population* will be used for either improving estimation efficiency or heterogeneity analysis. This application cannot be used to demonstrate empirically the other possible roles of covariate adjustment in RD designs, which nonetheless we also briefly discuss below. Finally, we offer practical guidance and briefly address other empirical and methodological uses of covariates in RD-based research. See, in particular, [15] for a recent methodological overview and further references. General-purpose software and replication files are available at <https://rdpackages.github.io/>.

Discussion of pros and cons

The RD design was introduced by [35] in 1960, but it only began to be recognized as a leading research design for observational studies forty years later: it was first formalized by [25] from a continuity-based (local regression) perspective, and later by [12] from a local-randomization (experimental) perspective. In empirical economics, the former conceptual perspective is often adopted because it justifies a local least squares regression analysis. While a wealth of methodological approaches are already available for the analysis and interpretation of RD designs, the role of covariates in particular remains an active area of methodological research. There is still no consensus on how to leverage covariates in RD applications among practitioners, a fact that has led to different valid (and invalid) empirical approaches, which can hamper the transparency and replicability of the empirical findings and policy recommendations.

Employing the Head Start application, we discuss the two main distinct roles of covariate adjustment in RD designs: for improving efficiency and for heterogeneity analysis. We also provide a unified framework to discipline practice leveraging covariates in empirical work. Moreover, we review other roles that covariates can take in RD applications conceptually, and caution against invalid practices sometimes found in the empirical literature.

Leveraging Covariates for Efficiency

Using the continuity-based framework, [9] first studied and formalized the role of covariate adjustment for efficiency purposes in RD designs. Their main goal was to provide an easy-to-interpret and easy-to-implement methodology to incorporate covariates in the otherwise canonical local linear regression approach for estimation of the (sharp) RD treatment effect [10]. Building on the standard analysis of experiments, the goal was not to change the parameter of interest but rather to improve the statistical precision in estimating the canonical RD average treatment effect at the cutoff.

We illustrate the main idea using the Head Start application. Figure 1(a) presents an RD Plot [11] employing a global linear regression, which suggests a sizable RD treatment effect. The first column in Table 1, labeled “canonical”, presents a modern implementation of the RD design originally studied in [30]. The estimated policy effect is a reduction of 2.41pp in child mortality among those U.S. counties that were offered Head Start, which is statistically significant at the 95% conventional level when using the robust bias-corrected statistical inference approach of [10]. The estimated RD treatment effect is highly statistically significant and represents a roughly 80% reduction in child mortality relative to the average mortality rate among control counties. This classical empirical finding is illustrated using a local RD Plot in Figure 1(b).

A key insight from [9] is that, to increase estimation efficiency, pre-intervention covariates should be incorporated with a restricted, common parameter across both control and treatment groups to obtain a consistent and more precise estimator of the canonical (sharp) RD treatment effect, under minimal identification assumptions and with low misspecification bias. A local regression model with full interactions between treatment assignment and pre-intervention covariates would require demeaning the covariates, thereby increasing the smoothing bias and hence reducing the precision of the RD estimator: for example, the corresponding mean square error (MSE) optimal bandwidth would be smaller than the MSE-optimal bandwidth proposed by [9], which in turn would lead to RD estimates with a larger variance on average.

The last two columns in Table 1 correspond to the approach proposed by [9] when either the eight covariates for efficiency are included in the estimation, or all nine pre-intervention covariates are incorporated in the analysis. The second column in Table 1, labeled “with covariates for efficiency”, reports an estimated policy effect corresponding to a reduction of 2.53pp in child mortality with an interval length reduction of about 3%, while the last column reports a reduction of 2.51pp in child mortality with a slightly larger interval length reduction. All these estimated effects are highly statistically significant, and quite similar across each other in magnitude. Notice that the models are nested: the first column does not include covariate adjustment, the next column adds eight covariates in the estimation procedure, and then the last column includes one additional covariate for efficiency purposes (*population*). The last column corresponds to the specification used in [9].

The Head Start application highlights some important empirical and methodological regularities for the role of covariate adjustment for efficiency purposes in RD designs. First, and perhaps

most importantly, the point estimates should not change from a statistical (and hence substantive) perspective. Second, the bandwidth employed should change to account for the additional covariates entering the estimation procedure. Third, the average length of the confidence intervals should decrease as more relevant pre-intervention covariates are included in the estimation. These implications are a consequence of properly employing pre-intervention covariates in modern RD estimation and inference. Importantly, it is always good practice to check that the covariates used are indeed unaffected by the policy: this is often done by producing RD plots and conventional RD estimates using the pre-intervention covariate as an outcome. For an example, see our companion replication files for the analog of Figure 1 and the first column in Table 1 using *population* as the outcome variable.

A limitation of leveraging covariates for efficiency is related to p-hacking, that is, an adversarial researcher could exploit the presence of many pre-intervention covariates to isolate the one local least squares regression specification that delivers statistically significant results, disregarding the fact that many regression models were fitted during the search process. This problem is not specific to RD designs and is fortunately mitigated by noting that the point estimates obtained from the canonical and the covariate-adjusted models should be similar to begin with. Thus, researchers should always report both estimates, canonical and covariate-adjusted, together with their accompanying robust bias-corrected statistical inference results. See Table 1 for one example.

The literature on covariate adjustment for efficiency improvements in RD designs has progressed rapidly in recent years. Two interesting, recent contributions are [4] and [31]. The first paper develops high-dimensional penalization methods to incorporate “many” covariates into RD estimation and inference, which can be used to remove the aforementioned concerns regarding model specification search (p-hacking). The second paper considers a local empirical likelihood approach, and also presents improved methodological results following [9].

Leveraging Covariates for Heterogeneity

The second main role of covariates in the RD design is for heterogeneity analysis. In this case, the pre-intervention covariates change the parameter of interest and thus play a different substantive role: the goal is to learn an RD treatment effect for each value of the covariates. As a result, estimation and inference methods need to be different from those discussed in the preceding section, which has led to the proliferation of applied approaches to leverage covariates for heterogeneity analysis in RD designs. Due to the curse of dimensionality, it is often the case that empirical work focuses on heterogeneity analysis one covariate at the time. Thus, the two most common empirical approaches are (i) to subset the data according to a discrete covariate, or otherwise discretized continuous covariate, and (ii) to incorporate a continuous covariate via joint local linear regression estimation with interactions between the treatment assignment variable and the covariate.

Perhaps surprisingly, there is only a handful of papers studying the methodological aspects of covariate adjustment for heterogeneity analysis in RD designs. [26] propose hypothesis testing

methods to determine whether there is heterogeneity underlying the canonical RD treatment effect, [34] and [2] discuss machine learning methods for estimating heterogeneous RD treatment effects when many continuous covariates are available, and [7] studies identification of weighted average treatment effects. None of these papers, however, develop estimation and inference methods based on the most common empirical practice for heterogeneity analysis in RD designs. [8] is the first paper to study the most common empirical practice based on estimation and inference via local least squares regressions with interactions.

We illustrate how to incorporate a covariate to conduct valid heterogeneity analysis using the Head Start data, following the methods recommended by [8]. To simplify the discussion, we only study heterogeneous RD treatment effects for small and large counties as determined by their population size: we discretize *population* into a binary variable indicating whether a county has at least 10k residents (henceforth, defined “large” counties), and then estimate the (conditional) RD treatment effects for each of the two subgroups with and without covariate adjustments for efficiency using the other eight pre-intervention covariates. Before starting with the empirical analysis, we test whether the probability of being a large county is continuous at the cutoff, an important empirical falsification test. We estimate the RD effect using the indicator for large counties as the outcome and find an estimated effect of 0.035 with a p -value of 0.704. The replication files give details, and also produce the corresponding RD Plot.

To begin the heterogeneity analysis, Figure 2(a) presents a global RD Plot for the two subgroups with linear regressions. The plot suggests a more sizable effect for the large counties when compared to small counties. The first two columns in Table 2 correspond to the heterogeneous RD analysis. The point estimates for the two subgroups are similar, but the RD treatment effect is only statistically significant for large counties. Specifically, it is found that Head Start decreases child mortality among highly populated counties by roughly 70% relative to the average mortality rate of highly populated control counties. Figure 2(b) gives a local RD Plot depicting these empirical findings.

The last two columns in Table 2 present a heterogeneous RD analysis leveraging the other pre-intervention covariates for efficiency. In this specific illustration, adding the other available covariates to improve precision does not help much. This is an empirical example where the additional covariates are likely to be uncorrelated with the treatment within each subgroup, thereby not offering noticeable efficiency improvements. Furthermore, in this application studying heterogeneous RD treatment effects leads to relatively small sample sizes, which calls for extra care in the analysis and interpretation of the results.

The Head Start application showcases one instance of heterogeneity analysis in RD designs where the findings indicate statistically significant policy effects for a particular subpopulation (largely populated counties). It also demonstrates an important limitation often encountered when estimating heterogeneous treatment effects: sample sizes are reduced because subpopulations are considered, hence leading to less precise policy estimates.

Leveraging Covariates for Other Purposes

Going beyond efficiency improvements and heterogeneity analysis, covariates have been used for other goals in RD designs. We overview these goals and provide further references to the literature. These methods cannot be illustrated using the Head Start application because they require specific data structure, or other departures beyond the canonical RD design setup, to be implemented.

Pre-intervention covariates can enhance RD designs by addressing issues like missing data, measurement error, and incorporating prior information via Bayesian methods. Measurement error in the RD running variable is a common concern, and recent approaches use covariates for correction [32, 5, 21, and references therein]. Under the assumption of missing-at-random, imputation methods can be applied directly within the local randomization framework. Bayesian methods also offer a distinct approach leveraging covariates to inform priors, or model distributions flexibly. For instance, prior experimental data has been used to inform Bayesian priors in RD studies [23], while principal stratification has been applied to estimate certain average treatment effects [20]. Although not yet widely adopted, these methods illustrate diverse roles for covariate adjustment in RD applications, primarily to refine estimation without altering the treatment effect parameter of interest.

Covariate adjustments in RD designs are also used to identify additional parameters of interest, such as those for certain subpopulations or for extrapolating treatment effects away from the cutoff. For example, in some RD designs, covariates define new subpopulations through varying cutoffs, as in Multi-Cutoff RD designs introduced by [16]. This approach explores heterogeneity by analyzing RD effects separately for each subgroup defined by its cutoff, which can also be aggregated through normalization and pooling under additional assumptions. These ideas are somewhat similar, but often conceptually distinct, from those related to heterogeneity analysis based on pre-intervention covariates as discussed previously. Similarly, [6] use treatment dosage changes across multiple cutoffs to explore counterfactual policies, and [17] define effects for subpopulations exposed to different cutoffs, both facilitating RD treatment effect extrapolation. Furthermore, in RD designs spanning multiple periods, treatment confounding can occur in one period but not another. The “differences in discontinuities” design of [24] estimates effects by comparing RD effects across periods, leveraging assumptions akin to parallel trends for unobservables across time. More recently, [27] studied dynamic RD designs where time plays an important role as the key additional covariate determining different treatment regimes.

When multiple cutoffs or periods are unavailable, extrapolation can still rely on auxiliary covariates. [3] introduce a conditional independence assumption that permits extrapolation across the running variable’s full support, while in geographic RD designs [29], assumptions focus on neighborhoods near geographic boundaries. For instance, [28] explore strategic location choices, imposing conditional independence within boundary neighborhoods (a local rather conditional independence assumption). Covariates also enhance extrapolation in other ways. [36] augment RD designs with exogenous outcome measures to estimate treatment effects beyond the cutoff under

suitable assumptions about outcome regression functions. See [17] for more discussion.

The references reviewed so far showcase a rich literature demonstrating how covariates in RD designs expand the scope of analysis by defining new parameters for subpopulations or treatment effects. These covariates may be inherent to the RD setup, such as multiple cutoffs or periods, or external, requiring separate collection, such as auxiliary characteristics or outcomes. Regardless, they can often be leveraged to improve scientific discoveries and evidence-based policy decisions.

Finally, it is sometimes claimed in the RD literature that covariates can be leveraged to “fix” a “broken” design. The rationale for adjusting covariate imbalances in RD designs draws on an analogy with flawed randomized experiments; settings where noticeable disparities in pre-intervention covariates arise, either due to defective randomization or random chance. However, RD designs present an important added complexity: lack of common support, which necessarily requires non-trivial extrapolation. This observation has prompted some researchers to employ covariate adjustment techniques to address shortcomings in the identification of canonical RD treatment effects. In the literature, RD analyses sometimes include fixed effects or other covariate adjustments, typically implemented as linearly additive regression terms in a (local) polynomial regression framework. Methodological advancements have proposed alternative approaches to manage covariate imbalances, such as multi-step non-parametric regression [22] and inverse probability weighting [33]. These techniques, developed within the continuity-based RD framework, aim to address cases where the conditional expectations of covariates given the running variable exhibit discontinuities near the cutoff, reflecting imbalances. As a result, they necessarily change the estimand.

Limitations and gaps

The RD design is a well-established methodology for program evaluation and causal inference, supported by numerous falsification tests and robustness checks to assess its validity [1, 18]. However, several important caveats should be noted. First, the policy effects identified by the RD design are “local” in nature, meaning they pertain only to units near the cutoff, rather than the entire population. Second, when pre-intervention covariates are included in an RD design, careful interpretation is required, and researchers must clearly articulate their intentions regarding these covariates. Third, examining heterogeneity in policy effects may significantly reduce the sample size available for analysis, which can affect statistical power. Finally, covariate adjustment cannot “fix” an RD design without making strong assumptions, and often changing the estimand.

Summary and policy advice

We examined the different roles that pre-treatment covariates can play in RD designs: enhancing the precision of estimates, exploring heterogeneity in treatment effects, and more broadly changing the estimand. We provided a detailed discussion of each of these roles and illustrated the first

two by revisiting a classic empirical study in the RD literature. Additionally, we emphasized the importance of clearly stating the purpose for including pre-treatment covariates in the analysis, as this clarifies the researcher's goal and facilitates the interpretation of empirical findings. Finally, we recommend always conducting empirical falsification testing to enhance the credibility of the findings, and also rely on well-developed methods for leveraging covariates in the analysis. Relaying on the empirical methods reviewed in this article will help replicability and comparability across studies.

Acknowledgments

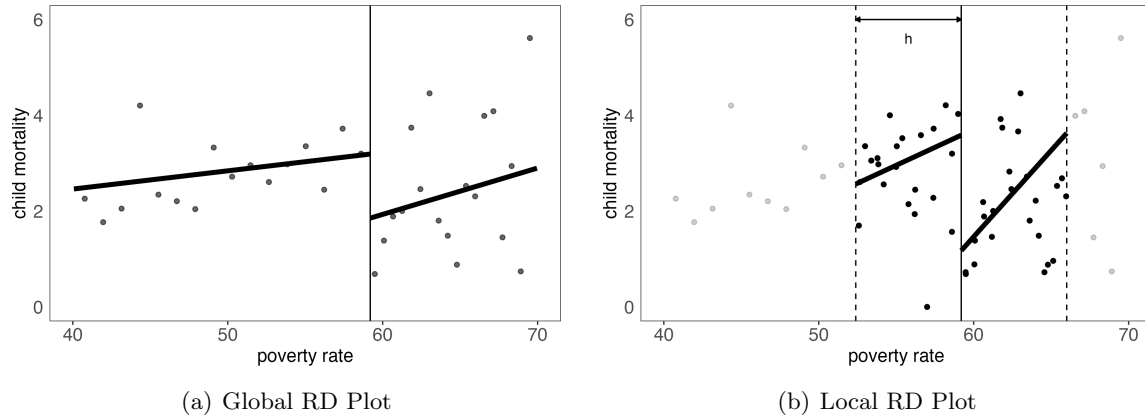
We thank Sebastian Calonico, Max Farrell, Luke Keele, and Rocio Titiunik for inspiring discussions. Cattaneo gratefully acknowledges financial support from the National Science Foundation through grants SES-2241575.

Competing interests

The IZA World of Labor project is committed to the [IZA Guiding Principles of Research Integrity](#). The authors declare that they have observed these principles.

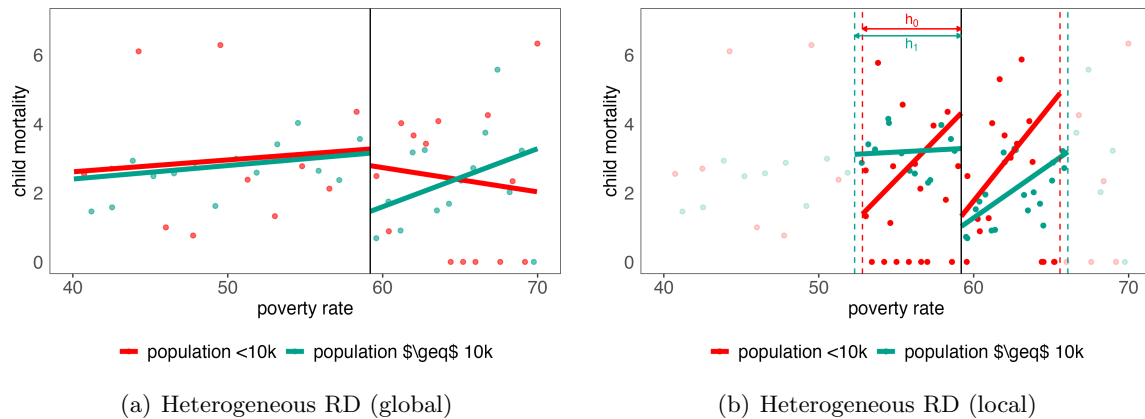
Figures

Figure 1: Head Start Canonical RD Average Treatment Effect



Notes: In all cases, we use a triangular kernel, employ a polynomial of degree 1 in the running variable, and equate the main bandwidth and the auxiliary bias bandwidth. No covariates are included for efficiency purposes. Dots represent averages within equally-spaced bins, solid lines represent separate linear fits on each side of the cutoff, and vertical dashed lines (when displayed) delimit the estimation sample of length twice the optimal MSE bandwidth. Sample sizes in panel (a) are $n_- = 2504$ and $n_+ = 300$, whereas in panel (b) $N_- = 234$ and $N_+ = 182$.

Figure 2: Head Start Heterogeneous RD Average Treatment Effects (*population*)



Notes: In all cases, we use a triangular kernel, use a polynomial of degree 1 in the running variable, and equate the main bandwidth and the auxiliary bias bandwidth. No covariates are included for efficiency. Dots represent averages within equally-spaced bins, solid lines represent separate linear fits on each side of the cutoff, and vertical dashed lines (when displayed) delimit the estimation sample of length twice the optimal MSE bandwidth. Sample sizes for small counties (*population* < 10k) in panel (a) are $n_- = 685$ and $n_+ = 101$, whereas in panel (b) $N_- = 70$ and $N_+ = 57$. Sample sizes for large counties (*population* $\geq 10k$) in panel (a) are $n_- = 1819$ and $n_+ = 199$, whereas in panel (b) $N_- = 162$ and $N_+ = 126$.

Tables

Table 1: Head Start Canonical RD Average Treatment Effect

type of RD	canonical	with covariates for efficiency	with all covariates
$\hat{\tau}$	-2.41	-2.53	-2.51
95% RCI	[-6.41, -1.09]	[-6.65, -1.5]	[-6.536, -1.4]
CI length change (%)	-	-3.25	-3.53
p -value	0.006	0.002	0.002
h	6.808	6.765	6.986
$N_- N_+$	234 180	231 179	242 184
% treatment effect	-78.56	-82.2	-81.5

Notes: In all cases, we use a triangular kernel, use a polynomial of degree 1 in the running variable, equate the estimation bandwidth and the bandwidth to estimate the bias, and use nearest neighbor heteroskedasticity-robust variance estimators. Column 2 contains results for the case in which no additional covariates are added to achieve efficiency gains; Column 3 shows the same statistics when 8 covariates are added (population between 14 and 17, population between 5 and 34, population with age 25+, % attending school between 14 and 17, % attending school between 5 and 34, % of urban areas, % of Black population); Column 4 also adds population as a covariate for efficiency gains. The rows report in order: the RD estimator; the bias-corrected 95% robust confidence interval; the percentage reduction in the RCI compared to the case without covariates; the p -value for the null that the true coefficient is 0; the optimal bandwidth; and the effective sample size to the left and right of the cutoff.

Table 2: Head Start Heterogeneous RD Average Treatment Effects

type of RD	canonical		with covariates for efficiency	
	<i>population</i> <10k	<i>population</i> \geq 10k	<i>population</i> <10k	<i>population</i> \geq 10k
$\hat{\tau}(z)$	-2.99	-2.26	-4.56	-2.48
95% RCI	[-11.49, 2.01]	[-6.1, -0.78]	[-9.56, 2.58]	[-6.65, -1.3]
CI length change (%)	-	-	-10.15	0.33
p -value	0.169	0.011	0.259	0.004
h	6.377	6.881	5.049	6.294
$N_- N_+$	70 55	162 126	56 45	145 116
% treatment effect	-103.11	-70.33	-157.88	-76.4

Notes: Panel A reports the results of the heterogeneity analysis obtained using two separate RD designs for the small and large counties. Panel B shows similar statistics using a single RD. Columns 2 and 3 contain results for the case where no additional covariates are added to achieve efficiency gains, whereas Panel (b) shows the same statistics when 8 covariates are added (population between 14 and 17, population between 5 and 34, population with age 25+, % attending school between 14 and 17, % attending school between 5 and 34, % of urban areas, % of Black population). Within each panel, the table shows in order: the RD estimator; the bias-corrected 95% robust confidence interval; the percentage reduction in the RCI compared to the case without covariates; the p -value for the null that the true coefficient is 0; the optimal bandwidth; and the effective sample size to the left and right of the cutoff. In all cases, we use a triangular kernel, use a polynomial of degree 1 in the running variable, and equate the estimation bandwidth and the bandwidth to estimate the bias.

References

- [1] Alberto Abadie and Matias D. Cattaneo. Econometric methods for program evaluation. *Annual Review of Economics*, 10:465–503, 2018.
- [2] Rafael Alcantara, Meijia Wang, P. Richard Hahn, and Hedibert Lopes. Modified bart for learning heterogeneous effects in regression discontinuity designs. *arXiv:2407.14365*, 2024.
- [3] Joshua D Angrist and Miikka Rokkanen. Wanna get away? regression discontinuity estimation of exam school effects away from the cutoff. *Journal of the American Statistical Association*, 110(512):1331–1344, 2015.
- [4] Yoichi Arai, Taisuke Otsu, and Myung Hwan Seo. Regression discontinuity design with potentially many covariates. *arXiv:2109.08351*, 2024.
- [5] Otavio Bartalotti, Quentin Brummet, and Steven Dieterle. A correction for regression discontinuity designs with group-specific mismeasurement of the running variable. *Journal of Business & Economic Statistics*, 39(3):833–848, 2021.
- [6] Marinho Bertanha. Regression discontinuity design with many thresholds. *Journal of Econometrics*, 218(1):216–241, 2020.
- [7] Carolina Caetano, Gregorio Caetano, and Juan Carlos Escanciano. Robust identification in regression discontinuity designs with covariates. *working paper*, 2024.
- [8] Sebastian Calonico, Matias D. Cattaneo, Max H. Farrell, Filippo Palomba, and Rocio Titiunik. Heterogeneous treatment effects in regression discontinuity designs. *working paper*, 2025.
- [9] Sebastian Calonico, Matias D. Cattaneo, Max H. Farrell, and Rocio Titiunik. Regression discontinuity designs using covariates. *Review of Economics and Statistics*, 101(3):442–451, 2019.
- [10] Sebastian Calonico, Matias D. Cattaneo, and Rocio Titiunik. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326, 2014.
- [11] Sebastian Calonico, Matias D. Cattaneo, and Rocio Titiunik. Optimal data-driven regression discontinuity plots. *Journal of the American Statistical Association*, 110(512):1753–1769, 2015.
- [12] Matias D. Cattaneo, Brigham Frandsen, and Rocio Titiunik. Randomization inference in the regression discontinuity design: An application to party advantages in the u.s. senate. *Journal of Causal Inference*, 3(1):1–24, 2015.
- [13] Matias D. Cattaneo, Nicolas Idrobo, and Rocio Titiunik. *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge Elements: Quantitative and Computational Methods for Social Science, Cambridge University Press, New York, NY, 2020.

- [14] Matias D. Cattaneo, Nicolas Idrobo, and Rocio Titiunik. *A Practical Introduction to Regression Discontinuity Designs: Extensions*. Cambridge Elements: Quantitative and Computational Methods for Social Science, Cambridge University Press, New York, NY, 2024.
- [15] Matias D. Cattaneo, Luke Keele, and Rocio Titiunik. Covariate adjustment in regression discontinuity designs. In D. S. Small J. R. Zubizarreta, E. A. Stuart and P. R. Rosenbaum, editors, *Handbook of Matching and Weighting in Causal Inference*, chapter 8, pages 153–168. Chapman & Hall, Boca Raton, FL, 2023.
- [16] Matias D. Cattaneo, Luke Keele, Rocio Titiunik, and Gonzalo Vazquez-Bare. Interpreting regression discontinuity designs with multiple cutoffs. *Journal of Politics*, 78(4):1229–1248, 2016.
- [17] Matias D. Cattaneo, Luke Keele, Rocio Titiunik, and Gonzalo Vazquez-Bare. Extrapolating treatment effects in multi-cutoff regression discontinuity designs. *Journal of the American Statistical Association*, 116(536):1941–1952, 2021.
- [18] Matias D. Cattaneo and Rocio Titiunik. Regression discontinuity designs. *Annual Review of Economics*, 14:821–851, 2022.
- [19] Matias D. Cattaneo, Rocio Titiunik, and Gonzalo Vazquez-Bare. Comparing inference approaches for rd designs: A reexamination of the effect of head start on child mortality. *Journal of Policy Analysis and Management*, 36(3):643–681, 2017.
- [20] Siddhartha Chib and Liana Jacobi. Bayesian fuzzy regression discontinuity analysis and returns to compulsory schooling. *Journal of Applied Econometrics*, 31(6):1026–1047, 2016.
- [21] Laurent Davezies and Thomas Le Barbanchon. Regression discontinuity design with continuous measurement error in the running variable. *Journal of econometrics*, 200(2):260–281, 2017.
- [22] Markus Frölich and Martin Huber. Including covariates in the regression discontinuity design. *Journal of Business & Economic Statistics*, 37(4):736–748, 2019.
- [23] Sara Geneletti, Aidan G O’Keeffe, Linda D Sharples, Sylvia Richardson, and Gianluca Baio. Bayesian regression discontinuity designs: Incorporating clinical knowledge in the causal analysis of primary care data. *Statistics in medicine*, 34(15):2334–2352, 2015.
- [24] Veronica Grembi, Tommaso Nannicini, and Ugo Troiano. Do fiscal rules matter? *American Economic Journal: Applied Economics*, 8(3):1–30, 2016.
- [25] Jinyong Hahn, Petra Todd, and Wilbert van der Klaauw. Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1):201–209, 2001.
- [26] Yu-Chin Hsu and Shu Shen. Testing treatment effect heterogeneity in regression discontinuity designs. *Journal of Econometrics*, 208(2):468–486, 2019.

- [27] Yu-Chin Hsu and Shu Shen. Dynamic regression discontinuity under treatment effect heterogeneity. *Quantitative Economics*, 15(4):1035–1064, 2024.
- [28] Luke Keele, Rocio Titiunik, and Jose R Zubizarreta. Enhancing a geographic regression discontinuity design through matching to estimate the effect of ballot initiatives on voter turnout. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 178(1):223–239, 2015.
- [29] Luke J Keele and Rocio Titiunik. Geographic boundaries as regression discontinuities. *Political Analysis*, 23(1):127–155, 2015.
- [30] Jens Ludwig and Douglas L. Miller. Does head start improve children’s life chances? evidence from a regression discontinuity design. *Quarterly Journal of Economics*, 122(1):159–208, 2007.
- [31] Jun Ma and Zhengfei Yu. Empirical likelihood covariate adjustment for regression discontinuity designs. *arXiv preprint arXiv:2008.09263*, 2024.
- [32] Zhuan Pei and Yi Shen. The devil is in the tails: Regression discontinuity design with measurement error in the assignment variable. In Matias D. Cattaneo and Juan Carlos Escanciano, editors, *Regression Discontinuity Designs: Theory and Applications (Advances in Econometrics, volume 38)*, pages 455–502. Emerald Group Publishing, 2017.
- [33] Sida Peng and Yang Ning. Regression discontinuity design under self-selection. In *International Conference on Artificial Intelligence and Statistics*, pages 118–126. PMLR, 2021.
- [34] Ágoston Reguly. Heterogeneous treatment effects in regression discontinuity designs. *arXiv preprint arXiv:2106.11640*, 2021.
- [35] Donald L. Thistlethwaite and Donald T. Campbell. Regression-discontinuity analysis: An alternative to the ex-post facto experiment. *Journal of Educational Psychology*, 51(6):309–317, 1960.
- [36] Coady Wing and Thomas D Cook. Strengthening the regression discontinuity design using additional design elements: A within-study comparison. *Journal of Policy Analysis and Management*, 32(4):853–877, 2013.